

Extremal Optimization at the Phase Transition of the 3-Coloring Problem

Stefan Boettcher*

Physics Department, Emory University, Atlanta, Georgia 30322, USA

Allon G. Percus†

*Computer & Computational Sciences Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA
and UCLA Institute for Pure and Applied Mathematics, Los Angeles, CA 90049, USA*

(Dated: February 2, 2008)

We investigate the phase transition of the 3-coloring problem on random graphs, using the extremal optimization heuristic. 3-coloring is among the hardest combinatorial optimization problems and is closely related to a 3-state anti-ferromagnetic Potts model. Like many other such optimization problems, it has been shown to exhibit a phase transition in its ground state behavior under variation of a system parameter: the graph's mean vertex degree. This phase transition is often associated with the instances of highest complexity. We use extremal optimization to measure the ground state cost and the “backbone”, an order parameter related to ground state overlap, averaged over a large number of instances near the transition for random graphs of size n up to 512. For graphs up to this size, benchmarks show that extremal optimization reaches ground states and explores a sufficient number of them to give the correct backbone value after about $O(n^{3.5})$ update steps. Finite size scaling gives a critical mean degree value $\alpha_c = 4.703(28)$. Furthermore, the exploration of the degenerate ground states indicates that the backbone order parameter, measuring the constrainedness of the problem, exhibits a first-order phase transition.

PACS numbers: 02.60.Pn, 05.10.-a 75.10.Nr,

I. INTRODUCTION

The most challenging instances of computational problems are often found near a critical threshold in the problem's parameter space [1], where certain characteristics of the problem change dramatically. One such problem, already discussed in Ref. [2, 3, 4, 5], is the 3-coloring problem. Consider a random graph [6] having n vertices and m edges placed randomly among all possible pairs of vertices. The number of edges emanating from each vertex is then Poisson-distributed around a mean degree $\alpha = 2m/n$. To 3-color the graph, we need to assign one of three colors to each vertex so as to minimize the number of “monochromatic” edges, i.e., those connecting vertices of the same color. In particular, we may want to decide whether it is possible to make an assignment without any monochromatic edges using, for instance, a backtracking assignment procedure. Typically, if the mean degree α is low (for example, when each vertex most likely has fewer than 3 neighbors), one quickly finds a perfect coloring. If the mean degree is high, one soon determines that monochromatic edges are unavoidable after fixing just a small number of vertices. At an intermediate mean degree value, however, some graphs are perfectly colorable while others are not. In that case, for each instance one must inspect many almost-complete colorings, most of which fail only when trying to assign a few final vertices, before colorability can be decided [4, 7]. For increasing

n , the regime of mean degree values α for which the decision problem is hard becomes narrowly focused, while the computational complexity of the backtracking algorithm within this regime grows faster than any power of n , signs of the impending singularity associated with a phase transition.

Such findings have spawned considerable interest among computer scientists and statistical physicists alike. On one hand, there appear to be close links to the properties of spin glass systems [8]. Using replica symmetry breaking, it was recently argued [2, 9] that 3-coloring undergoes a colorability transition at $\alpha_{\text{crit}} = 4.69$, heralded by the spontaneous emergence at $\alpha = 3.35$ of a sizable 3-core [10] that becomes over-constrained at the transition. This analysis shows, furthermore, that the hardest instances to decide are located between a clustering transition at $\alpha \approx 4.42$ and α_{crit} . On the other hand, attempts have been made to relate the nature of the phase transition to complexity classifications developed by computer scientists for combinatorial problems [11]: it has been suggested that NP-complete problems, which are hard to solve, display a first-order phase transition while easier problems lead to a second-order transition [12]. Such a relation, while intriguing and suggestive, is bound to be questionable in light of the fact that these phase transitions are based on a notion of average-case complexity (for instance, 3-coloring averaged over the ensemble of random graphs), distinct from the notion of worst-case complexity used by computer scientists to define NP-completeness [13] and other such complexity theoretic categories. In fact, it currently appears that a first-order transition is merely an indicator for the complexity of certain types of algorithms: local searches. A model prob-

*Electronic address: sboettc@emory.edu

†Electronic address: percus@ipam.ucla.edu

lem with a discontinuous transition, K -XORSAT [14, 15], can be solved by a fast global algorithm yet it is extremely hard for local search or backtracking assignment algorithms.

In this paper we consider the 3-coloring problem mentioned above. The problem is among the hardest combinatorial optimization problems, making it difficult to study the asymptotic properties of its phase transition with exact methods. It is also of considerable interest in its own right as a model for many practical optimization problems [11], and it is of some physical relevance due to its close relation to Potts anti-ferromagnets [16]. Some aspects of the 3-coloring phase transition have previously been explored [3, 4, 5]. In particular, Culberson and Gent [5] have studied the phase transition for random graphs with *exact* methods by “growing” random graphs of size $n \leq 225$, sequentially adding random edges to an existing graph to increase α , and checking along the way whether the graph is still colorable. Once a graph becomes uncolorable, it is discarded from the list of growing graphs, so the set of graphs becomes increasingly less representative of the ensemble when passing through the transition. In the process, these authors have evaluated the constrainedness of the variables in the graph, studying in detail many aspects of the approach to the transition. However, the main quantity they measure, called the “spine” [17], is in general an upper bound on the order parameter we consider here, since contributions from uncolorable graphs are neglected.

We investigate the properties near the phase transition by applying an optimization heuristic called *extremal optimization* (EO) [18, 19]. EO was recently introduced as a general-purpose optimization method based on the dynamics of driven, dissipative systems [20]. Our study shows that EO is capable of determining many ground-state properties efficiently, even at the phase transition [21]. EO performs a local neighborhood search that does not get stuck in local minima but proceeds to explore near-optimal configurations broadly. Hence, it is particularly well suited to measure global properties of the configuration space. Here, we use it to estimate the “backbone”, an overlap property between the highly degenerate ground state configurations that provides a more convenient order parameter than measuring mutual overlaps of all ground states [12]. While EO is not exact, benchmark comparisons with exactly-solved, large instances justify our confidence in its results. Our biggest errors originate from the lack of statistics at large n .

Our results indicate that the transition in the backbone size is of first order, though with only a small discontinuity. In fact, the discontinuity does not arise uniformly for all graphs in the ensemble, but is due to a fraction of instances that have a strong backbone of a characteristic size while the rest have hardly any backbone at all.

Using the procedure of Ref. [22] to control the quality of finite size scaling for the ground state cost function, we estimate the location of the transition as $\alpha_{\text{crit}} \approx 4.703(28)$, where the numbers in parentheses denote the

statistical error bar in the final digits. This is consistent with the presumably correct value of $\alpha = 4.69$ given by replica symmetry breaking methods [2] (see also Refs. [3, 4] for earlier estimates). We measure the size of the scaling window as $n^{-1/\nu}$ with $\nu = 1.43(6)$, close to the value of 1.5 estimated for 3-SAT [23], although it may be that trivial $n^{-1/2}$ -fluctuations from the variables not belonging to the 3-core [10] dominate at much larger n than considered here [24].

In the following section, we introduce the problem of 3-coloring in more detail and discuss the relevant observables we measure in order to analyze the phase transition. In Sec. III we discuss our EO implementation and its properties. In Sec. IV we present the results of our measurements, and we conclude with Sec. V.

II. 3-COLORING OF RANDOM GRAPHS

A random graph [6] is constructed from a set of n vertices by assigning an edge to $m = \alpha n/2$ of the $\binom{n}{2}$ pairs of vertices with equal probability, so that α is the average vertex degree. Here, we will only consider the regime of “sparse” random graphs where $m = O(n)$ and $\alpha = O(1)$. The goal of graph coloring is to label each vertex with a different color so as to avoid monochromatic edges.

Three different versions of the coloring problem are of interest. First, there is the classic problem of determining the “chromatic number” for a given graph, i.e., the minimum number of colors needed to color the graph while avoiding monochromatic edges. It is very difficult to devise a heuristic for this problem [25]. In the other two versions, we are given a fixed number K of colors to select from. The decision problem, K -COL, addresses the question of whether a given graph is colorable or not. Finally, the optimization problem, MAX- K -COL, tries to minimize the number of monochromatic edges (or equivalently, maximize the number of non-monochromatic edges, hence its name). Clearly, if we define the number of monochromatic edges as the “cost” or “energy” of a color assignment, determining whether the minimal cost is zero or non-zero corresponds to solving the decision problem K -COL, so finding the actual cost of the ground state is always at least as hard. Much of the discussion regarding the complexity near phase transitions in the computer science literature is focused on the decision problem [3, 5]. From a physics perspective, it seems more intuitive to examine the behavior of the ground states as one passes the transition. Accordingly, we will focus on the MAX- K -COL problem in this paper.

All these versions of coloring are NP-hard [11], and thus computationally hard in the worst case. To determine exact answers would almost certainly require a computational time growing faster than any power of n . Thus, extracting results about asymptotic properties of the problems is a daunting task, calling for the use of accurate heuristic methods, as discussed in the following

section.

The control parameter describing our ensemble of instances is the average vertex degree α of the random graphs. Constructing an appropriate order parameter to classify the transition is less obvious. The analogy to spin-glass theory [8, 26] suggests the following reasoning. In a homogeneous medium possessing a single pure equilibrium state, the magnetization provides the conjugate field to analyze the ferromagnetic transition. For our 3-coloring problem, the disorder induced by the random graphs leads to a decomposition into many coexisting but unrelated pure states with a distribution of magnetizations. Since the colors correspond to the spin orientations in the related Potts model, we need in principle to determine, for each graph, the overlap between all pairs of ground state colorings. Finally, this distribution has to be averaged over the ensemble. To simplify the task, one can instead extract directly the “backbone”, which is the set of variables that take on the same state in *all* possible ground state colorings of a given instance [12]. But even determining the backbone is a formidable undertaking: it requires not only finding a lowest cost coloring but sampling a substantial number of those colorings for each graph, since the ground state entropy is extensive.

Another level of difficulty arises due to the invariance of the ground states under a global color permutation. Thus, in the set of all ground states, each vertex can take on any color and the backbone as defined above is empty. To avoid this triviality, one may redefine the backbone in the following way. Instead of considering individual vertices, consider all *pairs* of vertices that are not connected by an edge [5]. Define the pair to be part of the *frozen* backbone if its vertices are of like color (monochromatic) in all ground state colorings, so that the presence of an edge there would necessarily incur a cost. Define the pair to be part of the *free* backbone if its vertices are of unlike color (non-monochromatic) in all ground state colorings, so that the presence of an edge there would not incur any cost. Since the fraction of pairs that belong to the frozen backbone measures the constrainedness of an instance, it is the relevant order parameter. We have also sampled the free backbone. As shown in Sec. IV, both seem to exhibit a first-order transition, though the jump for the frozen backbone is small.

By definition, the location of the transition is determined through a (second-order) singularity in the cost function C : the cost is asymptotically vanishing below the transition, it is continuous at the transition, and above it is always non-zero. We have therefore measured the ground state cost, averaged over many instances, for a range of mean degree values α and sizes n .

III. EXTREMAL OPTIMIZATION

To investigate the phase transition in 3-COL, we employ the extremal optimization heuristic (EO) [18]. The use of a heuristic method, while only approximate, allows

us to measure observables for much larger system sizes n and with better statistics than would be accessible with exact methods. We will argue below that we can obtain optimal results with sufficient probability that even systematic errors in the exploration of ground states will be small compared to the statistical sampling error.

Our EO implementation is as follows. Assume we are given a graph with a (however imperfect) initial assignment of colors to its vertices. Each vertex i has α_i edges to neighboring vertices, of which $0 \leq g_i \leq \alpha_i$ are “good” edges, i.e., to neighbors of a different color (not monochromatic). We define for each vertex a “fitness”

$$\lambda_i = \frac{g_i}{\alpha_i} \in [0, 1], \quad (1)$$

and determine a permutation Π (not necessarily unique) over the vertices such that

$$\lambda_{\Pi(1)} \leq \lambda_{\Pi(2)} \leq \dots \leq \lambda_{\Pi(n)}. \quad (2)$$

At each update step, EO draws a number k from a distribution

$$P(k) \sim k^{-\tau} \quad (1 \leq k \leq n) \quad (3)$$

with a bias toward small numbers. A vertex i is selected from the ordered list in Eq. (2) according to its “rank” k , i.e., $i = \Pi(k)$. Vertex i is updated *unconditionally*, i.e., it always receives a new color, selected at random from one of the other colors. As a consequence, vertex i and all its neighbors change their fitnesses λ and a new ranking Π will have to be established. Then, the update process starts over with selecting a new rank k , and so on until some termination condition is reached. Along the way, EO keeps track of the configuration with the best coloring it has visited so far, meaning the one that minimizes the number of monochromatic edges, $C = \sum_i (\alpha_i - g_i)/2$.

Previous studies have found that EO obtains near-optimal solutions for a variety of hard optimization problems [20] for a carefully selected value of τ [27, 28, 29]. For 3-COL, initial trials have determined that best results are obtained for the system sizes $n = 32, 64, \dots, 512$ at a (fixed) value of $\tau \approx 2.2$. This rather large value of τ [28] helps explore many low-cost configurations efficiently; if we merely wanted to determine low-cost solutions, larger values of n could have been reached more efficiently at a smaller value of τ .

It should be noted that our definition of fitness does not follow the generic choice $\lambda_i = g_i/2$ that would give a total configuration cost of $C = \text{const.} - \sum_i \lambda_i$. While this formulation sounds appealing, and does produce results of the same quality, our choice above produces those same results somewhat faster; there appears to be some advantage to treating all vertices, whose individual degrees α_i are Poisson-distributed around the mean α , on an equal footing. Furthermore, our implementation limits itself to partially sorting the fitnesses on a balanced heap [18], rather than ranking them perfectly as in Eq. (2). In this way, the computational cost is reduced by a factor of n while performance is only minimally affected [18].

A. Measuring the backbone

The backbone, described in Sec. II, is a collective property of degenerate ground states for a given graph. Thus, in this study we are interested in determining not only the cost C of the ground state, but also a good sampling of *all* possible ground state configurations. Local search with EO is ideally suited to probe for properties that are broadly distributed over the configuration space, since for small enough τ it does not get trapped in restricted regions. Even after EO has found a locally minimal cost configuration, it proceeds to explore the configuration space widely to find new configurations of the same or lower cost, as long as the process is run.

Against these advantages, one must recognize that EO is merely a heuristic approximation to a problem of exponential complexity. Thus, to safeguard the accuracy of our measurements, we devised the following adaptive procedure. For each graph, starting from random initial colorings, EO was run for n^3 update steps, using a minimum of 5 different restarts. For the lowest cost seen so far, EO keeps a buffer of up to $n/4$ most recently visited configurations with that cost. If it finds another configuration with the same cost, it quickly determines whether it is already in the buffer. If not, EO adds it on top of the buffer (possibly “forgetting” an older configuration at the bottom of the buffer). Thus, EO does not keep a memory of all minimal cost configurations seen so far, which for ground states can have degeneracies of $> 10^6$ even at $n = 64$. Instead of enumerating all ground states exhaustively, we proceed as follows. When EO finds a new, lowest cost configuration, it assumes initially that all pairs of equally colored vertices are part of the frozen backbone and all other pairs are part of the free backbone. If another configuration of the same cost is found and it is not already in the buffer, EO checks all of the pairs in it. If a pair has always been frozen (free) before and is so now, it remains part of the frozen (free) backbone. If a pair was always frozen (free) before and it is free (frozen) in this configuration, it is eliminated from both backbones. If a pair has already been eliminated previously, no action is taken. In this way, certain ground-state configurations may be missed or tested many times over, without affecting the backbones significantly.

Eventually, even if new and unrecognized configurations of the lowest cost are found, no further changes to either backbone are likely to occur. This fact motivates our adaptive stopping criterion for EO. Assume the current backbone was last modified in the r_0 th restart. Then, for this graph EO restarts for a total of at least $r = r_0 + \max\{r_0, 5\}$ times, terminating only when there has been no updates to the backbone over the previous $\max\{r_0, 5\}$ restarts. Of course, every time a new, lower-cost configuration is found, the buffer and backbone arrays are reset. Ultimately, this procedure leads to adaptive runtimes that depend on the peculiarities of each graph. The idea is that if the lowest cost state is found

in the first start and the backbone does not change over 5 more restarts, one assumes that no further changes to it will ever be found by EO. However, if EO keeps updating the backbone through, say, the 20th restart, one had better continue for 20 more restarts to be confident of convergence. The typical number of restarts was about 10, while for a few larger graphs, more than 50 restarts were required.

A majority of our computational time is spent merely confirming that the backbones have converged, since during the final $\max\{r_0, 5\}$ restarts nothing new is found. Nevertheless, EO still saves vast amounts of computer time and memory in comparison with exact enumeration techniques. The trade-off lies in the risk of missing some lowest cost configurations, as well as in the risk of never finding the true ground state to begin with. To get an estimate of the systematic error resulting from these uncontrollable risks, we have benchmarked our EO implementation against a number of different exact results.

First, we used a set of 700 explicitly 3-colorable graphs over 7 different sizes, $n = 75, 100, \dots, 225$ (100 graphs per value of n) at $\alpha = 4.7$, kindly provided by J. Culberson, for which exact spine values [17] were found as described in Ref. [5]. For colorable graphs such as these, the spine is identical to the backbone. Our EO implementation correctly determined the 3-colorability of all but one graph, and reproduced nearly all backbones exactly, regardless of size n . For all graph sizes, EO failed to locate colorable configurations on at most 5 graphs out of 100, and in those cases overestimated either backbone fraction by less than 4%. Only at $n = 225$ did EO miss the colorability of a single graph to find $C = 1$ instead, thereby underestimating both backbones.

In a different benchmark, containing colorable as well as uncolorable graphs, we generated 440 random graphs over 4 different sizes, $n = 32, 64, 128, 256$ and 11 different mean degree values $\alpha = 4.0, 4.1, \dots, 5.0$ (10 graphs per value of n and α). We found the exact minimum cost and exact fraction of pairs belonging to the backbone for these graphs, by removing edges until an exact branch-and-bound code due to M.A. Trick [30] determined 3-colorability. For example, finding that a graph had a ground state cost of $C = 2$ involved considering all possible 2-edge removals until a remainder graph was found to be 3-colorable. We then added edges to vertex pairs in this remainder graph, checking whether the graph stayed colorable: if so, that pair was eliminated from the frozen backbone. Likewise, we merged vertex pairs, checking whether the graph stayed colorable: if so, that pair was eliminated from the free backbone. Finally, we would repeat the procedure on all colorable 2-edge-removed remainder graphs, potentially eliminating pairs from the frozen and free backbones each time. Comparing with the exact testbed arising from this procedure shows that, for all graphs, EO found the correct ground state cost. Moreover, EO overestimated the frozen backbone fraction on only 4 graphs out of 440 (2 at $n = 128$ and 2 at $n = 256$, in both cases at $\alpha = 4.6$), and by

at most 0.004. This leads to a predicted systematic error that is at least an order of magnitude smaller than the statistical error bars in the results we present later. EO's free backbone results were not of as good quality, overestimating the backbone fraction on 36 graphs out of 440, by an average of 0.003 though in one case ($n = 256$, $\alpha = 4.6$) by as much as 0.027. The resulting systematic error, however, is still small compared to the statistical error bars in our main results.

It is also instructive to study the running times for EO, how they scale with increasing graph size, and how they compare with the exact algorithm we have used for benchmarking. In our EO implementation, we have measured the average number of update steps it took (1) to find the ground state cost for the first time, $\langle t_{\text{GS}} \rangle$, and (2) to sample the backbone completely, $\langle t_{\text{BB}} \rangle$. (Note that t_{BB} , corresponding to r_0 in Sec. III A, is always less than half the total time spent to satisfy the stopping criterion for an EO run described above.) Both t_{GS} and t_{BB} can fluctuate widely for graphs of a given n and α , especially when $C > 0$. However, since our numerical experiments involve a large number of graphs, the average times $\langle t_{\text{GS}} \rangle$ and $\langle t_{\text{BB}} \rangle$ are reasonably stable. Furthermore, $\langle t_{\text{GS}} \rangle$ and $\langle t_{\text{BB}} \rangle$ show only a weak dependence on α , varying by no more than a factor of 2: $\langle t_{\text{GS}} \rangle$ increases slowly for increasing α , and $\langle t_{\text{BB}} \rangle$ has a soft peak at α_c [31]. Thus, for each n we average these times over α as well, leading only to a slight increase in the error bars (on a logarithmic scale). We have plotted the average quantities in Fig. 1 on a log-log scale as a function of n . It suggests that the time EO takes to find ground states increases exponentially but very weakly so. Once they are found, EO manages to sample a sufficient number of them in about $O(n^{3.5})$ updates to measure the backbone accurately.

By contrast, one cannot easily quantify the scaling behavior of running times for the exact branch-and-bound benchmarking method. As ground state cost increases, the complexity of the method quickly becomes overwhelming, and rules out using it to measure average quantities with any statistical significance. Clearly, branch-and-bound itself has exponential complexity for determining colorability. For the sizes studied here, however, the exponential growth in n appears in fact subdominant to the $O(n^{C+2})$ complexity of evaluating the backbone for a graph with non-zero ground state costs. When $C = 1$ or 2, the combinatorial effort is manageable, but at $n = 256$, graphs just at the transition ($\alpha = 4.70$) reach $C \geq 3$ and the algorithm takes weeks to test all remainder graphs. From this comparison, one can appreciate EO's speed in estimating the backbone fractions, however approximate!

IV. NUMERICAL RESULTS

With the EO implementation as described above, we have sampled ground state approximations for a large number of graphs at each size n . In particular we have




FIG. 1: Log-log plot of the average time $\langle t_{\text{GS}} \rangle$ to reach the lowest cost state found (circles) and $\langle t_{\text{BB}} \rangle$ to sample the backbone (squares), in units of EO-update steps, as a function of system size n . The dashed line, $0.003n^{3.5} \exp(0.004n)$ gives a reasonable fit to $\langle t_{\text{GS}} \rangle$ (after having assumed the $n^{3.5}$ power law for the fit), and the dash-dotted line, $0.3n^{3.5}$, is a fit obtained for $\langle t_{\text{BB}} \rangle$. (Taking these crude fits at face value, merely reaching the first good ground state approximations would begin to dominate the runtime at about $n \approx 10^3$).

considered, over a range of α , 100 000 random graphs of size $n = 32$, 10 000 of size $n = 64$, 4 000 of size $n = 128$, and 1 000 of size $n = 256$. By averaging over the lowest energies found for these graphs, we obtain an approximation for average ground state costs $\langle C \rangle$ as a function of α and n , as shown in Fig. 2. We have also sampled 160 instances of size $n = 512$, which provided enough statistics for the backbone though not for the ground state costs.

With a finite size scaling ansatz

$$\langle C \rangle \sim n^\delta f[(\alpha - \alpha_{\text{crit}})n^{1/\nu}], \quad (4)$$

systematically applied [22], it is possible to extract precise estimates for the location of the transition α_c and the scaling window exponent ν . In the scaling regime, one might assume that the cost for the fixed argument of the scaling function is independent of the size, i.e., $\delta = 0$, indicated by the fact that for all values of n the cost functions cross in virtually the same point. Hence, in results we have previously reported [19], we obtained what appeared to be the best data collapse by fixing $\delta = 0$ and choosing $\alpha_{\text{crit}} = 4.72(1)$ and $\nu = 1.53(5)$, with the error bars in parentheses being estimates based on our own assessment of the data collapse. But a more careful automated fit to our data, provided to us by S.M. Bhattacharjee, gives $\delta \approx -0.001(3)$, $\alpha_{\text{crit}} = 4.703(28)$, and $\nu = 1.43(6)$ with a tolerance level of $\eta = 1\%$ (see Ref. [22]). While these fits are consistent with our previous results, they are also consistent with and much closer to the presumably exact result of $\alpha_{\text{crit}} = 4.69 \dots$ [2], and the error estimates are considerably more trustworthy.

The scaling window is determined by two competing contributions: for the intermediate values of n accessi-




FIG. 2: Plot of the average cost as a function of the vertex degree α . After correct finite size scaling, the data collapses onto a single scaling function, as shown in the insert. The fit gives $\alpha_{\text{crit}} \approx 4.70$, marked by a vertical line.

ble in this study it is dominated by nontrivial contributions arising from the correlations amongst the variables, which yields $\nu \approx 1.43(6)$, similar to satisfiability problems [23]. However, for sufficiently large n , Wilson [24] has shown that $\nu \geq 2$, due to intrinsic features of the ensemble of random graphs. The argument may be paraphrased as follows. Since $\alpha = O(1)$ and vertex degrees are Poisson-distributed with mean α , a finite fraction of vertices in a random graph have degrees 0, 1, or 2 (those not belonging to the 3-core [10]) and thus cannot possibly cause monochromatic edges. But this finite fraction itself undergoes (normal) $\sim 1/\sqrt{n}$ fluctuations, and these fluctuations limit the narrowing of the cost function's scaling window at large n . Such variables make up about 15% of the total near α_{crit} , so we estimate the crossover to occur at $n^{-1/\nu} \sim 0.15n^{-1/2}$ or $n \approx 10^5$, assuming all other constants to be unity.

Our next main result is the estimate of the backbone near the phase transition, as described in Sec. III A. We have sampled the frozen and the free backbones [5] separately. Our results show the fraction of vertex pairs in each backbone, and are plotted in Fig. 3. For the free backbone, consistent with our definition, we do not include any pairs that are already connected by an edge. Although they make up only $O(1/n)$ of the pairs, the inclusion of these would cause a significant finite size effect when the backbone is small, and only by omitting them does the free backbone vanish for $\alpha < \alpha_{\text{crit}}$. In principle, according to our definition one should also be sure to exempt from the frozen backbone any pairs that are connected by a monochromatic edge in all ground state configurations, but at α_{crit} their impact on the backbone is only $O(1/n^2)$.

As Fig. 3 shows, both backbones appear to evolve toward a discontinuity for increasing n . The backbone fraction comes increasingly close to vanishing below α_{crit} , followed by an increasingly steep jump and then a plateau




FIG. 3: Plot of the frozen (top) and the free (bottom) backbone fraction as a function of the vertex degree α . The critical point $\alpha_{\text{crit}} \approx 4.70$ is indicated by a vertical line.

that, to within statistical noise, appears stable at large n . The height of the plateau at α_{crit} suggests that on average about 6% of all pairs are frozen and close to 20% are free, with both values rising further for increasing degree. The “jump” in the frozen backbone is somewhat smaller than that in the free backbone, adding a higher degree of uncertainty to that interpretation, although still well justified within the error bars. Indeed, given the considerable ground state degeneracies, it would be surprising if the frozen backbone were large.

A more detailed look at the data (Fig. 4) suggests that the distribution of frozen backbone fractions for individual instances is bimodal at the transition, i.e., about half of the graphs have a backbone well over 10% while the other half have no backbone at all, leading to the average of 6% mentioned above. Furthermore, there appears to be some interesting structure in the backbone discontinuity, which may be significant beyond the noise. Note in Fig. 3 that for larger n , the increase of the frozen backbone stalls or even reverses right after the jump before rising further. This property coincides with the emergence of non-zero costs in the ground state colorings (see Fig. 2). The sudden appearance of monochromatic edges seems initially to reduce the frozen backbone fraction:

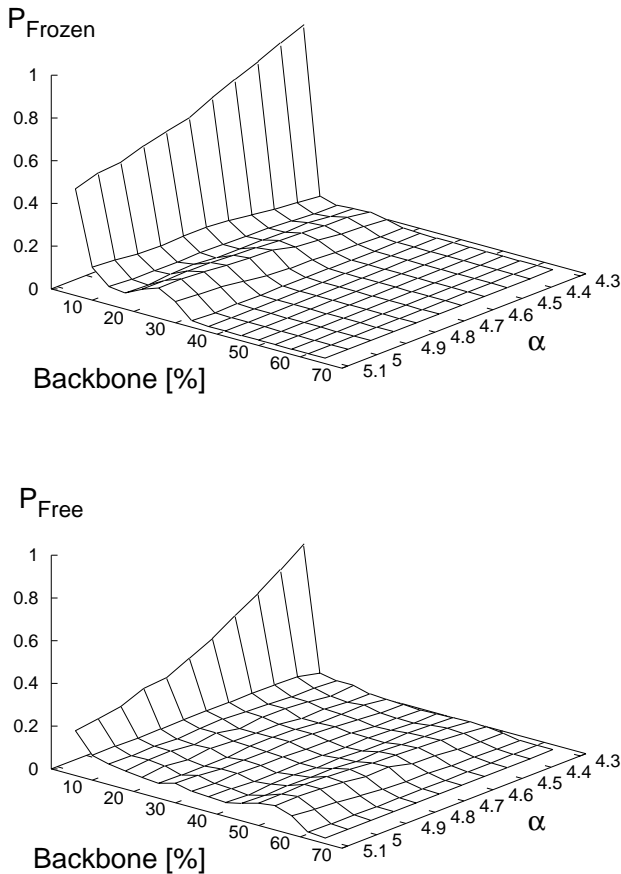


FIG. 4: Plot of the typical frozen (top) and free (bottom) backbone probability, obtained here from the $n = 128$ graphs. For mean degree values $\alpha < \alpha_{\text{crit}}$, graphs are almost certain to have a vanishing backbone, both frozen and free. Above α_{crit} a majority of graphs still do not exhibit any backbone, but a finite fraction of graphs display a sizable backbone fraction, clustered at a characteristic size. The average backbones plotted in Figs. 3 represent the average of these apparently bimodal distributions. This qualitative picture appears to hold for increasing n , although $n > 128$ data are somewhat noisy.

typically there are numerous ways of placing those few edges, often affecting the most constrained variables pairs and eliminating them from the frozen backbone. Similar observations have been made by Culberson [5]. According to this argument, only the frozen backbone should exhibit such a stall (or dip). Indeed, Fig. 3 indeed shows a less hindered increase in the free backbone, though the difference there may be purely due to statistical noise.

V. CONCLUSIONS

We have considered the phase transition of the MAX-3-COL problem for a large number of instances of random

graphs, of sizes up to $n = 512$ and over a range of mean degree values α near the critical threshold. For each instance, we have determined the fraction of vertex pairs in the frozen and free backbones, using an optimization heuristic called *extremal optimization* (EO) [18]. Based on previous studies [21], EO is expected to yield an excellent approximation for the cost and the backbone. Comparisons with a testbed of exactly-solved instances suggest that EO's systematic error is negligible compared to the statistical error.

Using a systematic procedure for optimizing the data collapse in finite size scaling [22], we have argued that the transition occurs at $\alpha_{\text{crit}} = 4.703(28)$, consistent with earlier results [3, 4, 5, 19] as well as with a recent replica symmetry breaking calculation yielding 4.69 [2]. We have also studied both free and frozen backbone fractions around the critical region. A simple argument [12] demonstrates that below the critical point the backbone fraction always vanishes for large n . At and above the critical point, neither backbone appears to vanish, suggesting a first-order phase transition. This is in close resemblance to K -SAT for $K = 3$ [12]; indeed, both are computationally hard at the threshold.

Even though the backbone is defined in terms of minimum-cost solutions, its behavior appears to correlate more closely with the complexity of finding a zero-cost solution (solving the associated decision problem) at the threshold. One possible explanation is that instances there have low cost, so finding the minimal cost is only polynomially more difficult than determining whether a zero-cost solution exists. Interestingly, our 3-coloring backbone results mirror those found for the spine [5], an upper bound on the backbone that is defined purely with respect to zero-cost graphs. The authors of that study speculate that at the threshold, although the spine is discontinuous, the backbone itself might be *continuous*. Our results contradict this speculation, instead providing support for a relation — albeit restricted — between backbone behavior and average-case complexity.

Acknowledgments

We are greatly indebted to Somen Bhattacharjee for providing an automated evaluation of our finite size scaling fit. We wish to thank Gabriel Istrate, Michelangelo Grigni, and Joe Culberson for helpful discussions, and the Los Alamos National Laboratory LDRD program for support.

-
- [1] B.A. Huberman and T. Hogg, *Artificial Intelligence* **33**, 155–171 (1987).
 - [2] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, *Phys. Rev. Lett.* **89**, 268701 (2002).
 - [3] P. Cheeseman, B. Kanefsky, and W.M. Taylor, in *Proc. of IJCAI-91*, eds. J. Mylopoulos and R. Rediter (Morgan Kaufmann, San Mateo, CA, 1991), 331–337.
 - [4] T. Hogg, *Artificial Intelligence* **81**, 127–154 (1996).
 - [5] J. Culberson and I.P. Gent, *J. Theor. Comp. Sci.* **265**, 227–264 (2001).
 - [6] B. Bollobas, *Random Graphs*, (Academic Press, London, 1985).
 - [7] S. Cocco and R. Monasson, *Phys. Rev. Lett.* **86**, 1654–1657 (2001), and *Euro. Phys. J. B* **22**, 505–531 (2001).
 - [8] M. Mezard, G. Parisi, and M.A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
 - [9] A. Braunstein, R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, *Phys. Rev. E* **68**, 036702 (2003).
 - [10] B. Pittel, J. Spencer, and N. Wormald, *J. Comb. Th. B* **67**, 111 (1996).
 - [11] M.R. Garey and D.S. Johnson, *Computers and Intractability, A Guide to the Theory of NP-Completeness* (W.H. Freeman, New York, 1979).
 - [12] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky, *Nature* **400**, 133–137 (1999), and *Random Struct. Alg* **15**, 414–435 (1999).
 - [13] S.A. Cook, *Proc. 3rd Annual ACM Symp. on Theory of Computing*, 151–158 (Assoc. for Computing Machinery, New York, 1971).
 - [14] F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, *Phys. Rev. E* **63**, 026702 (2001).
 - [15] M. Leone, F. Ricci-Tersenghi, and R. Zecchina, *J. Phys. A Math. Gen.* **34**, 4615–4626 (2001).
 - [16] F.Y. Wu, *Rev. Mod. Phys.* **54**, 235–268 (1982); R. Kotecky, *Phys. Rev. B* **31**, 3088 (1985); P. Svenson and M.G. Nordahl *Phys. Rev. E* **59**, 3983 (1999).
 - [17] B. Bollobas, C. Borgs, J.T. Chayes, J.H. Kim, and D.B. Wilson, *Scaling Window of the 2-SAT Transition*, *Random Structures and Algorithms* **18**, 201–256 (2001).
 - [18] S. Boettcher and A.G. Percus, *Artificial Intelligence* **119**, 275–286 (2000).
 - [19] S. Boettcher and A.G. Percus, *Phys. Rev. Lett.* **86**, 5211–5214 (2001).
 - [20] S. Boettcher, A.G. Percus, and M. Grigni, *Lecture Notes in Computer Science* **1917**, 447–456 (2000).
 - [21] S. Boettcher, *J. Phys. A: Math. Gen.* **32**, 5201–5211 (1999).
 - [22] S.M. Bhattacharjee and F. Seno *J. Phys. A: Math. Gen.* **34**, 6375 (2001).
 - [23] S. Kirkpatrick and B. Selman, *Science* **264**, 1297–1301 (1994).
 - [24] D.B. Wilson, *The empirical values of the critical k-SAT exponents are wrong*, preprint.
 - [25] D.S. Johnson, C.R. Aragon, L.A. McGeoch, and C. Schevon, *Operations Research* **39**, 378–406 (1991).
 - [26] K.H. Fischer and J.A. Hertz, *Spin Glasses* (Cambridge University Press, Cambridge, 1991).
 - [27] S. Boettcher and A.G. Percus, *Phys. Rev. E* **64**, 026114 (2001).
 - [28] S. Boettcher and M. Grigni, *J. Phys. A: Math. Gen.* **35**, 1109–1123 (2002).
 - [29] J. Dall, *Searching Complex State Spaces with Extremal Optimization and other Stochastics Techniques*, Master Thesis, Fysisk Institut, Syddansk Universitet Odense, 2000 (in danish).
 - [30] A. Mehotra and M.A. Trick, *A column generation approach to graph coloring*, Technical Report Series, Graduate School of Industrial Administration, Carnegie Mellon University, 1995. Graph coloring code available at <http://mat.gsia.cmu.edu/COLOR/solvers/trick.c>.
 - [31] While search times to solve the decision problems should be sharply peaked here, determining the lowest cost solutions remains hard when $C > 0$ for $\alpha > \alpha_c$, leading more to a plateau than to a peak. It may also be the case that a search with EO is less influenced by the transition, as other studies [21] have suggested, but the range around α_c that we have studied here is too small to be conclusive on this question.